

Terms used

Found 90,429 of 166,357

[storage](#) [first](#) [property](#) [record](#) [data](#) [second](#) [access](#) [data](#) [third](#) [spreadsheet](#) [file](#) [index](#) [cell](#)

Sort results by

relevance

Display results

expanded form


[Save results to a Binder](#)

[Search Tips](#)

☐ Open results in a new window

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 


1 [Fortran 8X draft](#)



Loren P. Meissner

December 1989 **ACM SIGPLAN Fortran Forum**, Volume 8 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(21.36 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)


Standard Programming Language Fortran. This standard specifies the form and establishes the interpretation of programs expressed in the Fortran language. It consists of the specification of the language Fortran. No subsets are specified in this standard. The previous standard, commonly known as "FORTRAN 77", is entirely contained within this standard, known as "Fortran 8x". Therefore, any standard-conforming FORTRAN 77 program is standard conforming under this standard. New features can b ...

2 [Fast detection of communication patterns in distributed executions](#)

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available:  [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...


3 [Elements of Data Management Systems](#)



George G. Dodd

June 1969 **ACM Computing Surveys (CSUR)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.22 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

4 [Archival storage for digital libraries](#)



Arturo Crespo, Hector Garcia-Molina

May 1998 **Proceedings of the third ACM conference on Digital libraries**

Publisher: ACM Press

Full text available:  [pdf\(1.32 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 [Special issue: AI in engineering](#)



Publisher: ACM Press

Full text available: pdf(8.79 MB) Additional Information: [full citation](#), [abstract](#)

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

6 Mobile data management: Mimic: raw activity shipping for file synchronization in mobile file systems



Tae-Young Chang, Aravind Velayutham, Raghupathy Sivakumar

June 2004 **Proceedings of the 2nd international conference on Mobile systems, applications, and services MobiSys '04**

Publisher: ACM Press

Full text available: pdf(334.54 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we consider the problem of file synchronization when a mobile host shares files with a backbone file server in a network file system. Several *diff* schemes have been proposed to improve upon the transfer overheads of conventional file synchronization approaches which use full file transfer. These schemes compute the binary *diff* of the new file with respect to the old copy at the server and transfer the computed *diff* to the server for file-synchronization. However ...

Keywords: file synchronization, mobile file system, raw activity shipping

7 Third Generation Computer Systems



Peter J. Denning

December 1971 **ACM Computing Surveys (CSUR)**, Volume 3 Issue 4

Publisher: ACM Press

Full text available: pdf(3.52 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The common features of third generation operating systems are surveyed from a general view, with emphasis on the common abstractions that constitute at least the basis for a "theory" of operating systems. Properties of specific systems are not discussed except where examples are useful. The technical aspects of issues and concepts are stressed, the nontechnical aspects mentioned only briefly. A perfunctory knowledge of third generation systems is presumed.

8 Experiments in social data mining: The TopicShop system



Brian Amento, Loren Terveen, Will Hill, Deborah Hix, Robert Schulman

March 2003 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 10 Issue 1

Publisher: ACM Press

Full text available: pdf(377.92 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Social data mining systems enable people to share opinions and benefit from each other's experience. They do this by mining and redistributing information from computational records of social activity such as Usenet messages, system usage history, citations, or hyperlinks. Some general questions for evaluating such systems are: (1) is the extracted information valuable? and (2) do interfaces based on the information improve user task performance? We report here on *TopicShop*, a system ...

Keywords: Cocitation analysis, collaborative filtering, computer-supported cooperative work, information visualization, social filtering, social network analysis


9 Pen computing: a technology overview and a vision



André Meyer

July 1995 **ACM SIGCHI Bulletin**, Volume 27 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(5.14 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This work gives an overview of a new technology that is attracting growing interest in public as well as in the computer industry itself. The visible difference from other technologies is in the use of a pen or pencil as the primary means of interaction between a user and a machine, picking up the familiar pen and paper interface metaphor. From this follows a set of consequences that will be analyzed and put into context with other emerging technologies and visions. Starting with a short historic ...


10 The Quadtree and Related Hierarchical Data Structures



Hanan Samet

June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(4.87 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 Designing and mining multi-terabyte astronomy archives: the Sloan Digital Sky



Survey

Alexander S. Szalay, Peter Z. Kunszt, Ani Thakar, Jim Gray, Don Slutz, Robert J. Brunner

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00**, Volume 29 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(429.09 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The next-generation astronomy digital archives will cover most of the sky at fine resolution in many wavelengths, from X-rays, through ultraviolet, optical, and infrared. The archives will be stored at diverse geographical locations. One of the first of these projects, the Sloan Digital Sky Survey (SDSS) is creating a 5-wavelength catalog over 10,000 square degrees of the sky (see <http://www.sdss.org/>). The 200 million objects in the multi-terabyte database will have mostly numerical attributes ...

Keywords: Internet, archive, astronomy, data analysis, data mining, database, scalable


12 Multidimensional access methods



Volker Gaede, Oliver Günther

June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.05 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Search operations in databases require special support at the physical level. This is true for conventional databases as well as spatial databases, where typical search operations include the point query (find all objects that contain a given search point) and the region query (find all objects that overlap a given search region). More than ten years of spatial database research have resulted in a great variety of multidimensional access methods to support ...

Keywords: data structures, multidimensional access methods


13 Speeding up construction of PMR quadtree-based spatial indexes



Gisli R. Hjaltason, Hanan Samet

October 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(355.72 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Spatial indexes, such as those based on the quadtree, are important in spatial databases for efficient execution of queries involving spatial constraints, especially when the queries involve spatial joins. In this paper we present a number of techniques for speeding up the

construction of quadtree-based spatial indexes, specifically the PMR quadtree, which can index arbitrary spatial data. We assume a quadtree implementation using the "linear quadtree", a disk-resident representation ...

Keywords: Bulk-loading, I/O, Spatial indexing

14 Comparison of access methods for time-evolving data



Betty Salzberg, Vassilis J. Tsotras

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Publisher: ACM Press

Full text available: [pdf\(529.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. The comparison is based on a collection of important performance criteria, including the space consumed, update processing, and query time for representative queries. The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequencies are made. When a number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

Keywords: I/O performance, access methods, structures, temporal databases

15 Session summaries from the 17th symposium on operating systems principle (SOSP'99)



Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Publisher: ACM Press

Full text available: [pdf\(3.15 MB\)](#) Additional Information: [full citation](#), [index terms](#)

16 Searching in high-dimensional spaces: Index structures for improving the performance of multimedia databases



Christian Böhm, Stefan Berchtold, Daniel A. Keim

September 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 3

Publisher: ACM Press

Full text available: [pdf\(1.39 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research issue in the field of multimedia databases is the content-based retrieval of similar multimedia objects such as images, text, and videos. However, in contrast to searching data in a relational database, a content-based retrieval requires the search of similar objects as a basic functionality of the database system ...

Keywords: Index structures, indexing high-dimensional data, multimedia databases, similarity search

17 Optimistic replication



Yasushi Saito, Marc Shapiro

March 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 1

Publisher: ACM Press

Full text available: [pdf\(656.72 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes

popular. Optimistic replication deploys algorithms not seen in tradition ...

Keywords: Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques


18 Query evaluation techniques for large databases



Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(9.37 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality


19 Personal distributed computing: the Alto and Ethernet software



Butler Lampson

January 1986 **Proceedings of the ACM Conference on The history of personal workstations**

Publisher: ACM Press

Full text available:  [pdf\(3.00 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



The personal distributed computing system based on the Alto and the Ethernet was a major effort to make computers help people to think and communicate. The paper describes the complex and diverse collection of software that was built to pursue this goal, ranging from operating systems, programming environments, and communications software to printing and file servers, user interfaces, and applications such as editors, illustrators, and mail systems.

20 Computing curricula 2001



September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press




Full text available:  [pdf\(613.63 KB\)](#)
 [html\(2.78 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)